Heart rate-based Identification of Users of IoT Wearables: A Supervised Learning Approach

Sachit A. J. Desa[1], Basem Suleiman[1,2], and Waheeb Yaqub[1]

[1] School of Computer Science, The University of Sydney, Australia
[2] School of Computer Science and Engineering, The University of New South Wales, Sydney, Australia
{sachit.desa, basem.suleiman, waheeb.faizmohammad}@sydney.edu.au

Abstract. Biometric identification from heart rate sequences provides a simple yet effective mechanism, that can neither be reverse engineered nor replicated, to protect user privacy. This study employs a highly efficient time series classification (TSC) algorithm, miniROCKET, to identify users by their heart rate. The approach adopted in this study employs user heart rate data, a simplified form of heart activity, captured during exercise, filtered, and contextualized within exercise routines, for user classification. Results from this study are empirically evaluated on a real-world data set, containing 115,082 workouts across 304 users, by three other state-of-the-art TSC algorithms. Our experiments showed that for 36 users, the variance explained by heart rate feature is 74.0%, when coupled with speed and altitude, the variance increases to 94.0%. For 304 users, the variance explained by heart-rate is 32.8% and increased to 65.9% with contextual features. This exploratory study highlights the potential of heart rate as a biometric identifier. It also underscores how contextual factors, such as speed and altitude change, can improve classification of timeseries data when coupled with smart data preprocessing.

1 Introduction

The popularity of Internet of Things (IoT) devices has been driven by their ability to record biometric and contextual data for use in providing personalised healthcare services. However, this popularity increases the risk of personal data associated with these devices being hacked - a risk magnified by the projected uptake of wearables, forecast to pass six billion by 2025 [1]. Tied to data collected by wearables, and IoT devices in their various forms, is the personal information of individuals that use wearables. Literature surveyed within this area highlighted two critical themes associated with IoT device data: a.) The metadata associated with wearables may be compromised in several instances due to currently employed device protection mechanisms [6], and b.) The limited amount of research, till date, in the privacy field of wearable devices [14].

The sparse amount of work completed in this domain has established how electrocardiogram (ECG) data, can be used for identification [3]. ECG signals have been showcased to identify individuals with an accuracy upwards of 95%

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_41

[15]. However, such data is collected within controlled environments and with specialist equipment. The novelty of this study is its use of heart rate data collected by commercial IoT devices, as opposed to specialist equipment (i.e. ECGs).

Currently, IoT devices call for user meta-data to be entered (name, date of birth, etc.) for easy identification. A large number of these devices offer a connection to phone and email services, as well as being a portal to other applications that store user data. Access to this information via passwords, fingerprints, or two-factor authentication can provide hackers with a trove of personal data, as well as information linked to institutions and third parties.

In this paper, we address the challenge of maintaining user privacy by proposing an approach that identifies users via their heart rate, instead of traditional authentication methods that are prone to hacking. The robustness of the proposed approach relies on the difficulty of replicating and reverse-engineering the heart-rate sequence. The goal of this paper is to contribute to the growing adoption of wearables (and personal IoT devices such as smart watches) whilst protecting user privacy. We introduce a new research direction in which we present the findings of an approach that employs key features known to affect heart rate (speed and altitude), and a filtering strategy, to identify users with their wearable devices. The contributions of this study are:

- Exploration of heart rate, collected from IoT wearables, as a unique identifying biometric feature.
- Filtering methods that allow for the most efficient and effective use of biometric data to identify users of wearable devices.
- An empirical study of the performance of time series classification (TSC) algorithms on heart rate data. We consider a real-world data set consisting of 115,082 workouts across 304 users of IoT devices.

2 Related Work

Access to personal information from wearable IoT devices and trusted third parties has been highlighted by several studies [6,12,14,16]. Specifically, activity pattern recognition [12], data driven privacy-setting recommendations [14], and data inference risks from third parties [16], are methods through which privacy breaches occur. Reichherzer et al. (2017) [12] present machine learning techniques that intercept, track, and classify individuals and their fitness activities.

The energy availability for wearable and battery-powered devices is limited, precluding the use of many sophisticated and established authentication mechanisms [2]. Similarly, PINs and biometrics are not suitable due to the limitations of small form factor [12]; biometric authentication such as face, fingerprint, and iris recognition can be counterfeited as they do not check for liveness in a subject [8]. The use of ECG data overcomes this limitation as it is an inherent measure of liveness. Studies have shown that whilst the duration of a heartbeat has been known to vary with stress, anxiety, and time of day, the structure of a heartbeat contains scalar differences to variations in stress [8]. Existing research,

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_41

including studies conducted by Israel et al. [9], Irvine et al. [7, 8], and Biel et al. [3], demonstrate the heartbeat structure to be unique to an individual.

For a set of 20 subjects, Biel et al. [3] use a Soft Independent Modeling of Class Analogy (SIMA) classifier to identify individuals by ECG data. SIMA returned a classification accuracy of 98.0%. In a separate study, Shen et al. [15] found, via feature extraction, that measurements of palm ECG signals on 168 individuals returned an accuracy of 95.3%. A key observation from ECG data was that signals from individuals with a similar age, weight, height, and gender can be significantly different [15]. Whilst ECG data contains multiple dimensions for classification, heart rate is simply a measure of heartbeats per minute. The overarching challenge identified as part of this study, is the use of a simplified measure of heart activity, heart rate, as an encoding mechanism for protecting user privacy. In lieu of data rich signals captured by ECG tests, completed when an individual is at rest, this study uses contextual data, recorded by wearables during exercise, for classification. In using heart rate data, collected by wearables, for user identification, this study aims to bridge the gap between current research that has focused exclusively on ECG data, collected in controlled environments with specialist equipment, and 'real-world' data that is coarse.

3 Datasets for Classification

The data used within this study, originally made available by Endomondo, comprised of 253,020 workouts across 1,104 users. Ni et al. [11] used the Endomondo data to generate a consolidated dataset containing users with a minimum of 10 workouts, restricted features to those relevant to this study, and interpolated time series features to obtain data at regular intervals. As part of this study, features such as latitude, longitude, and timestamp were re-imported, and other redundant features discarded.

Data sub-setting and partitioning: An initial exploration of the data found that the number of workouts per user varied significantly (i.e. from 10 to 1301 workouts per user) with the distribution skewed towards users with a lower number of workouts.

In view of achieving a more uniform distribution of workouts per user, reducing the target variable space, and offering sufficient data for each algorithm to be trained, data was subset by the number of records per user. The minimum number of workouts per user was based on a wide ranging study by Ruiz et al. (2021) [13] who completed a multi-variate time series classification study of 18 TSC algorithms over 30 datasets. The largest target variable space consisted of 39 classes, each associated with $\sim$ 170 records. An accuracy of 36.7% was achieved on this dataset. Dataset 1 for this study was filtered to contain 304 users with a minimum of 200 workout records.

Dataset 1 was subset to evaluate and compare the selected algorithms over a successively smaller number of classes and metrics. Subsetting was based on empirical evidence collected from classification tests on the University of California, Riverside (UCR) database. This database contains two heartbeat and three ECG

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_41

datasets with a maximum of 5 classes. The records per class range from 100 to 1000, with a maximum classification accuracy achieved for a dataset containing 442 records per class. To test the classification sensitivity associated with the number of workout records per user, three datasets were generated, containing a minimum of 350 and a maximum of 560 workout records per user. Each dataset was a subset of its parent, with the test and train partitions of dataset 1 carrying through to dataset 4. The properties of each dataset are presented in Table 1.

Table 1: Dataset used for the current study

Dataset	Records	Min. Records[1]	Users	Size (GB)
1	115,082	200	304	5.9
2	69,111	350	129	3.5
3	59,694	400	104	3.1
4	27,064	560	36	1.4

[1] Denotes the minimum number of exercise records per user

4 Methodology - Heart rate based Identification

The network architecture for this study consists of a filtering module, the mini-ROCKET algorithm, and a feature selection module. The filtering module pre-processes user data to determine the optimal number of activities, shortest date range, and least number of records required to generate a physiological profile. The optimal filtering parameters are applied to the test data to generate a filtered dataset. The miniROCKET algorithm completes classification by contextualising heart-rate with each feature (distance, altitude, etc..) and feature combination. The optimum feature combination is then retained for classification on the test data. For any dataset, the network determines the optimal number of records, shortest date range, and fewest activities needed for classification. The model also ranks and quantifies the contribution of each contextual feature to classification accuracy. Figure 1 presents an overview of the model architecture.

miniROCKET and comparative TSC algorithms: miniROCKET was used as the mainstay algorithm for this study. Ruiz et al. [13], who completed a wide-ranging study on multi-variate timeseries data, found **R**andom **C**onvolutional **KE**rnel **T**ransform (ROCKET) to outperform current time series classification algorithms in accuracy and computational efficiency.

ROCKET employs random convolutional kernels in conjunction with a linear classifier (either ridge regression or logistic regression) for classification. Kernels can be configured based on 5 different parameters: bias, size (length), weights, dilation, and padding. MiniROCKET, an updated version of ROCKET, minimizes the hyperparameter options for each kernel; four of the five hyperparameters are fixed, whilst the weights are selected as either -1 or 2. Fixing hyperameter values ensures that miniROCKET is considerably faster [4].

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_41

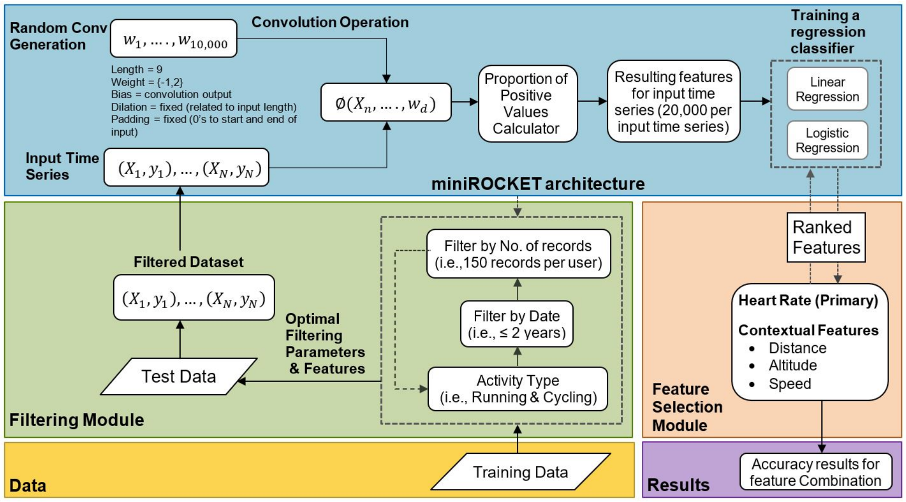

Fig. 1: Model architecture implemented for classification

Time Series Forest (TSF), InceptionTime, and Attention-Long Short-Term Memory (A-LSTM), were employed on a comparative basis. The two deep learning algorithms, Inception Time and A-LSTM, were selected for their accuracy and scalability [10], both of which were key themes underlying this study. TSF was selected for its emphasis on temporal characteristics, which would support the classification of heart-rate sequences, particularly when considering the temporal dependence of heart rate on speed and altitude [5].

Hyperparameter tuning and selection: Hyperparameter selection was completed on the test sets of datasets 1, 2, 3, and 4. For deep learning algorithms, InceptionTime and A-LSTM, the batch size and epochs were tuned. The number of epochs had a significant bearing on accuracy and computational time unlike batch size. The dilation parameter was tuned for miniROCKET, and number of estimators were optimized for TSF. For the parameter range over which accuracy remained consistent, hyperparameters were fine-tuned to optimise computational efficiency. Table 2 presents the final parameter range selected for the four algorithms across datasets.

Table 2: Final selected hyper-parameter values

Algorithm	Parameter	Value	Criteria
miniROCKET	Dilations	32	Selected for optimized computational time
InceptionTime	Batch size	32	Values selected for improved accuracy
	Epochs	180-220	
A-LSTM	Batch size	32	Influential over computational time
	Epochs	400	Influential over accuracy
TSF	Estimators	50	Fine-tuned to optimize computational time

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_41

5 Results and Discussion

Table 3 presents the final results, across datasets and feature combinations. The system architecture remained consistent for each of the selected algorithms. It is to be noted that because speed and time were included, the addition of distance was redundant. However, the distance feature resulted in an improved accuracy for IT, and has therefore been included in the results.

Table 3: Classiication accuracies across datasets 1 to 4

	Dataset 1				Dataset 2			
Feature List	**TSF**	**mRCKT**	**IT**	**ALSTM**	**TSF**	**mRCKT**	**IT**	**ALSTM**
HR, Speed, Alt. & Dist.	63.8	**65.9**	63.2	64.6	75.0	**78.8**	69.5	71.6
HR, Speed, & Altitude	**65.0**	64.8	53.8	61.8	**73.0**	72.2	48.2	68.6
HR & Speed	43.8	49.1	29.8	**58.7**	54.6	**60.0**	25.1	45.1
HR	25.1	32.8	7.6	**40.8**	33.9	45.6	12.7	**50.8**
	Dataset 3				**Dataset 4**			
Feature List	**TSF**	**mRCKT**	**IT**	**ALSTM**	**TSF**	**mRCKT**	**IT**	**ALSTM**
HR, Speed, Alt. & Dist.	78.4	**82.8**	67.5	78.0	86.7	**94.0**	77.4	89.1
HR, Speed, & Altitude	76.9	**81.2**	59.2	71.3	86.9	**93.4**	59.6	82.6
HR & Speed	57.6	**68.0**	40.0	66.3	72.5	**85.8**	38.2	80.5
HR	36.5	**49.2**	13.7	41.0	54.7	**74.0**	15.5	65.0

miniROCKET outperformed the baseline models for accuracy across all datasets containing the full set of features. With the exception of IT, other algorithms generated accuracies within 7.3% (±3.6) for datasets with all the features. Based on heart rate alone, which explains upto 40.8% of variance across 306 users, A-LSTM outperformed miniROCKET across datasets 1 and 2. The 'long-short term memory' component of A-LSTM is more effective in sequence detection across a large number of users than the linear classifier employed by miniROCKET. The accuracy of A-LSTM on heart activity is reinforced by Karim et al. [10] who record an accuracy of 95.0% for 5 classes on ECG data.

General trends indicate that accuracy is highly dependent on the number of users. For dataset 1, containing 304 users, the classification accuracy, with the exception of IT, is within 2.1%. The accuracy for dataset 4, containing 36 users, is within 7.3%. These results, presented in Figure 2, are indicative of the inability of algorithms to scale up with an increased number of user classes.

Trends observed in Figure 2 demonstrate accuracy to be linked with the number of users to be classified. Some linearity is observed in the results. It is noteworthy that there is a slight increase in accuracy recorded by IT between 104 and 129 users. The increase in accuracy, despite the increase in users, is reflective of the documented instability of the Inception Network as a result of the variability from random weight initialization. The current study aimed to

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_41

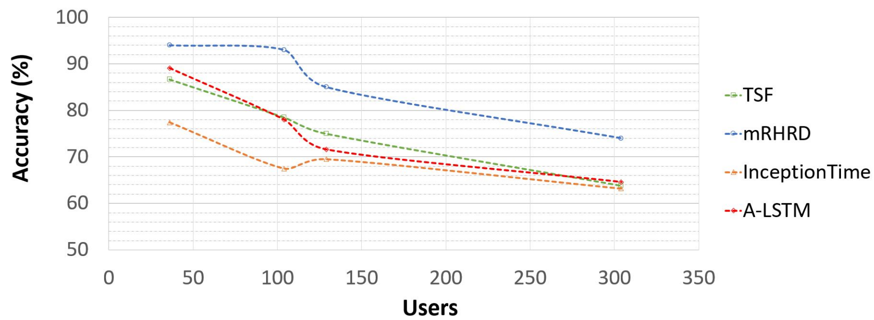

Fig. 2: Accuracy according to the number of users for all features

overcome this instability by implementing an ensemble of five Inception Networks, with every prediction attributed the same weight, however, this did not completely address the inherent algorithmic instability.

The results of our experimental procedure demonstrate that heart rate explains up to 40.8% of variance across 304 users, the variance explained increases to 66.0% when all features are considered. Similarly, the variance across 36 users is close to 74.0% when heart rate is considered, and increases to 94.0% when all features are accounted for. On average, the addition of speed increased classification accuracy by 15.1% (±3.3) and altitude increased accuracy by 12.3% (±3.2). Both speed and altitude are factors known to influence heart rate [5], in the present study their inclusion allowed for an improved classification accuracy.

Classification studies completed using ECG data, for 20 and 168 individuals, recorded accuracies of 98.0% and 95.3% respectively [3,15]. The dimensionality of ECG data, in the form of P, QRS complex, and T waves, is a critical differentiator between users, and device data. The sophistication of currently available wearables, at a consumer level, has not yet evolved to capture underlying heart rate patterns. This is a major challenge that prohibits the use of heart rate data as a bio-metric identifier when a large number of user classes is considered. The addition of speed and altitude to heart rate is more encouraging. With consumer grade wearables unable to collect data that is as rich as ECG signals, the addition of contextual data, and smart preprocessing, can improve user identification.

6 Conclusion and Future Work

This study demonstrates how heart rate, contextualised within within exercise routines and coupled with a smart filtering strategy, can be used as a biometric identifier. Our results establish that heart rate alone can be used to explain up to 40% of variance for 304 users and 74% of variance for 36 users. Whilst the variance explained by heart rate may not sufficiently distinguish between a large number of users, the addition of speed and altitude increases classification accuracy to 66% for 304 users and 94% for 36 users. Future work should consider the inclusion of features representing latent physiological information (i.e. calorie

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_41

count and hydration) to improve classification. By highlighting the importance of contextual data and smart filtering, this work establishes a strong baseline for researchers to develop a fully operational real-world classification model.

References

1. A survey on wearable technology: History, state-of-the-art and current challenges. Computer Networks **193**, 108074 (2021). `https://doi.org/https://doi.org/10.1016/j.comnet.2021.108074`
2. Bellovin, S.M., Merritt, M.: Encrypted key exchange: Password-based protocols secure against dictionary attacks. In: Proceedings 1992 IEEE Computer Society Symposium on Research in Security and Privacy (1992)
3. Biel, L., Pettersson, O., Philipson, L., Wide, P.: Ecg analysis: A new approach in human identification. Trans. on Instrumentation and Measure. pp. 808–12 (2001)
4. Dempster, A., Schmidt, D.F., Webb, G.I.: minirocket. In: Singapore (ed.) Proc. of the 27th ACM SIGKDD Conf. on Knowledge Discovery and Data Mining (2021)
5. Fleckenstein, D., Ueberschar, O., Wustenfeld, J.C., Rudrich, P., Wolfarth, B.: Effect of uphill running on vo2, heart rate and lactate accumulation on lower body positive pressure treadmills. Sports **9**, 4 (2021)
6. ul Haq, M.E., Malik, M., Azam, M., Naeem, U., Khalid, A., Ghazanfar, M.: Identifying users with wearable sensors based on activity patterns. In: The 11th Int. Conf. on Emerging Ubiquitous Systems and Pervasive Networks. Madeira (2020)
7. Irvine, J., Wiederhold, B., Gavshon, L., Israel, S., McGehee, S., Meyer, R., Wiederhold, M.: Heart rate variability: a new biometric for human identification. In: Vegas, L. (ed.) International Conference on Artificial Intelligence (IC-AI'2001) (2001)
8. Irvine, J.M., Israel, S.A., Scruggs, W.T., Worek, W.J.: eigenpulse: Robust human identification from cardiovascular function. Pattern Recognition pp. 3427–35 (2008)
9. Israel, S.A., Irvine, J.M., Cheng, A., Wiederhold, M.D., Wiederhold, B.K.: Ecg to identify individuals. Pattern Recognition **38**(1), 133–142 (2005)
10. Karim, F., Majumdar, S., Darabi, H., Chen, S.: Lstm fully convolutional networks for time series classification. CoRR **abs/1709.05206** (2017), `http://arxiv.org/abs/1709.05206`
11. Ni, J., Muhlstein, L., JulianMcAuley: "modeling heart rate and activity data for personalized fitness recommendation,". WWW '19: Proceedings of the 2019 World Wide Web Conference, San Francisco (2019)
12. Reichherzer, T., Timm, M., Earley, N., Reyes, N., Kumar, V.: Using machine learning techniques to track individuals their fitness activities. In: Proceedings of 32nd International Conference on Computers and Their Applications (2017)
13. Ruiz, A., Flynn, M., Large, J., Middlehurst, M., Bagnall, A.: The great multivariate time series classification bake off: a review and experimental evaluation of recent algorithmic advances. Data Mining and Knowledge Discovery **35**, 401–49 (2021)
14. Sanchez, O., Torre, I., He, Y., Knijnenburg, B.: A recommendation approach for user privacy preferences in the fitness domain. User Modeling and User-Adapted Interaction **30**, 513–65 (2019)
15. Shen, T.W., Tompkins, W.J., Hu, Y.H.: Implementation of a one-lead ecg human identification system on a normal population. Journal of Engineering and Computer Innovations pp. 12–21 (2011)
16. Torre, I., Koceva, F., Sanchez, O., Adorni, G.: Fitness trackers and wearable devices: How to prevent inference risks? In: 11th Int. Conf. on Body Area Ntwrks (2017)

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_41